



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 103 952 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
08.06.2005 Bulletin 2005/23

(51) Int Cl.7: **G10L 15/06**

(21) Application number: **00310492.4**

(22) Date of filing: **27.11.2000**

(54) Context-dependent acoustic models for speech recognition with eigenvoice training

Kontextabhängige akustische Modelle für die Spracherkennung mit Eigenstimmenanpassung

Modèles acoustiques contextuels pour la reconnaissance de parole avec un entraînement sur
vecteurs propres

(84) Designated Contracting States:
DE ES FR GB IT

(30) Priority: **29.11.1999 US 450392**

(43) Date of publication of application:
30.05.2001 Bulletin 2001/22

(73) Proprietor: **MATSUSHITA ELECTRIC INDUSTRIAL
CO., LTD.
Kadoma-shi, Osaka 571-8501 (JP)**

(72) Inventors:

- Kuhn, Roland
Santa Barbara, California 93110 (US)
- Contolini, Matteo
Santa Barbara, California 93109 (US)
- Junqua, Jean Claude
Santa Barbara, California 93110 (US)

(74) Representative: **Atkinson, Ralph et al
Atkinson Burrington,
25-29 President Buildings,
President Way
Sheffield S4 7UR (GB)**

(56) References cited:
EP-A- 0 953 968

- ROBERT WESTWOOD: "Speaker Adaptation Using Eigenvoices" 31 August 1999 (1999-08-31), WOLFSON COLLEGE, DEPARTMENT OF ENGINEERING, CAMBRIDGE UNIVERSITY, CAMBRIDGE, UK XP002176018 * abstract ** page 17-25 ** page 29-30 *
- WILLETT D ET AL: "REFINING TREE-BASED STATE CLUSTERING BY MEANS OF FORMAL CONCEPT ANALYSIS, BALANCED DECISION TREES AND AUTOMATICALLY GENERATED MODEL-SETS" PHOENIX, AZ, MARCH 15 - 19, 1999, NEW YORK, NY: IEEE, US, 15 March 1999 (1999-03-15), pages 565-568, XP000900183 ISBN: 0-7803-5042-1
- KUHN R ET AL: "FAST SPEAKER ADAPTATION USING A PRIORI KNOWLEDGE" PHOENIX, AZ, MARCH 15 - 19, 1999, NEW YORK, NY: IEEE, US, 15 March 1999 (1999-03-15), pages 749-752, XP000900229 ISBN: 0-7803-5042-1

Description

[0001] The present invention relates to a method for developing context-dependent models for automatic speech recognition.

5 [0002] Small vocabulary speech recognition systems have as their basic units the words in the small vocabulary to be recognized. For instance, a system for recognizing the English alphabet will typically have twenty-six models, one model per letter of the alphabet. This approach is impractical for medium and large vocabulary speech recognition systems. These larger systems typically take as their basic units the phonemes or syllables of a language. If a system contains one model (for example a Hidden Markov Model) per phoneme of a language, it is called a system with 10 "context-independent" acoustic models.

[0003] If a system employs different models for a given phoneme, depending on the identity of the surrounding phonemes, the system is said to employ "context-dependent" acoustic models. An allophone is a specialised version of a phoneme defined by its context. For instance, all the instances of "ae" pronounced before "t", as in "bat", "fat", etc define an allophone of "ae".

15 [0004] For most languages, the acoustic realization of a phoneme depends very strongly on the preceding and following phonemes. For example, an "eh" preceded by a "y" (as in "yes") is quite different from an "eh" preceded by "s" (as in "set"). Thus, for a system with a medium-sized or large vocabulary, the performance of context-dependent acoustic models is much better than that of context-independent models. Most practical applications of medium and large vocabulary recognition systems employ context-dependent acoustic models today.

20 [0005] Many context-dependent recognition systems today employ decision tree clustering to define the context-dependent, speaker-independent acoustic models. A tree-growing algorithm finds questions about the phonemes surrounding the phoneme of interest and splits apart acoustically dissimilar examples of the phoneme of interest. The result is a decision tree of yes-no questions for selecting the acoustic model that will best recognize a given allophone. Typically, the yes-no questions pertain to how the allophone appears in context (ie what are its neighbouring phonemes).

25 [0006] A conventional decision tree defines for each phoneme a binary tree containing yes/no questions in the root node and in each intermediate node (children, grandchildren, etc of the root node). The terminal nodes, or leaf nodes, contain the acoustic models designed for particular allophones of the phoneme. Thus, in use, the recognition system traverses the tree, branching "yes" or "no" based on the context of the phoneme in question until the leaf node containing the application model is identified. Thereafter the identified model is used for recognition.

30 [0007] Unfortunately, conventional allophone modelling can go wrong. The applicants believe that this is because current models do not take into account the particular idiosyncrasies of each training speaker. Current methods assume that individual speaker idiosyncrasies will be averaged out if a large pool of training speakers is used. However, in practice, it has been found that this assumption does not always hold. Conventional decision tree-based allophone models work fairly well when a new speaker's speech happens to resemble the speech of the training speaker population.

35 [0008] The thesis of Robert Westwood entitled "Speaker Adaptation using Eigenvoices" of 31 August 1999, in connection with work being done through Wolfson College, Department of Engineering, Cambridge University, Cambridge, United Kingdom XP002176018 describes a technique for the adaptation of speaker models. Examples are shown and

40 a conclusion is reached to the effect that for the Eigenvoice decomposition technique disclosed to be effective, it is important to model the interspeaker variability well.

[0009] According to an aspect of the present invention, there is provided a method of the aforesaid type, comprising generating an eigenspace to represent a training speaker population; providing a set of acoustic data for at least one training speaker and representing said acoustic data in said eigenspace to determine at least one allophone centroid for said training speaker; subtracting said centroid from said acoustic data to generate speaker-adjusted acoustic data for said training speaker; and using said speaker-adjusted acoustic data to grow at least one decision tree having leaf nodes containing context-dependent models for different allophones.

45 [0010] The invention will now be described by way of example only, with reference to the accompanying drawings, of which:

50 Figure 1 is a diagrammatic illustration of speaker space useful in understanding how the centroid of a speaker population and the associated allophone vectors differ from speaker to speaker;

Figure 2 is a block diagram of a first presently preferred embodiment called the eigencentroid plus delta tree embodiment;

55 Figure 3 illustrates one embodiment of a speech recognizer that utilises the delta decision tree developed from the embodiment illustrated in Figure 2;

Figure 4 is another embodiment of a speech recognizer that also uses the delta decision trees generated by the embodiment of Figure 2;

Figure 5 illustrates how a delta tree might be constructed using the speaker-adjusted data generated by the embodiment of Figure 2;

Figure 6 shows the grouping of speaker-adjusted data in acoustic space corresponding to the delta tree of Figure 5;

Figure 7 illustrates an exemplary delta decision tree that includes questions about the eigenspace dimensions; and

Figure 8 illustrates a second embodiment of the invention, useful in applications where there is more complete quantity of data per speaker.

[0011] The techniques of the invention may be applied to a variety of different speech recognition problems. The techniques are perhaps most useful in medium and large vocabulary applications, where it is not feasible to represent each full word by its own model. Two embodiments of the invention will be described here. It will be understood that the principles of the invention may be extended to other embodiments, as well.

[0012] The first embodiment is optimized for applications where each training speaker has supplied a moderate amount of training data: for example, on the order of twenty to thirty minutes of training data per speaker. With this quantity of training data it is expected that there will be enough acoustic speech examples to construct reasonably good context independent, speaker dependent models for each speaker. If desired, speaker adaptation techniques can be used to generate sufficient data for training the context independent models. Although it is not necessary to have a full set of examples of all allophones for each speaker, the data should reflect the most important allophones for each phoneme somewhere in the data (i.e., the allophones have been pronounced a number of times by at least a small number of speakers).

[0013] The recognition system of this embodiment employs decision trees for identifying the appropriate model for each allophone, based on the context of that allophone (based on its neighboring phonemes, for example). However, unlike conventional decision tree-based modeling systems, this embodiment uses speaker-adjusted training data in the construction of the decision trees. The speaker adjusting process, in effect, removes the particular idiosyncrasies of each training speaker's speech so that better allophone models can be generated. Then, when the recognition system is used, a similar adjustment is made to the speech of the new speaker, whereby the speaker-adjusted allophone models may be accessed to perform high quality, context dependent recognition.

[0014] An important component of the recognition system of this embodiment is the Eigenvoice technique by which the training speaker's speech, and the new speaker's speech, may be rapidly analyzed to extract individual speaker idiosyncrasies. The Eigenvoice technique, discussed more fully below, defines a reduced dimensionality Eigenspace that collectively represents the training speaker population. When the new speaker speaks during recognition, his or her speech is rapidly placed or projected into the Eigenspace to very quickly determine how that speaker's speech "centroid" falls in speaker space relative to the training speakers.

[0015] As will be fully explained, the new speaker's centroid (and also each training speaker's centroid) is defined by how, on average, each speaker utters the phonemes of the system. For convenience, one can think of the centroid vector as consisting of the concatenated Gaussian mean vectors for each state of each phoneme HMM in a context independent model for a given speaker. However, the concept of "centroid" is scalable and it depends on how much data is available per training speaker. For instance, if there is enough training data to train a somewhat richer speaker dependent model for each speaker (such as a diphone model), then the centroid for each training speaker could be the concatenated Gaussian means from this speaker dependent diphone model. Of course, other models such as triphone models and the like, may also be implemented.

[0016] Figure 1 illustrates the concept of the centroids by showing diagrammatically how six different training speakers A-F may pronounce phoneme 'ae' in different contexts. Figure 1 illustrates a speaker space that is diagrammatically shown for convenience as a two-dimensional space in which each speaker's centroid lies in the two-dimensional space at the center of the allophone vectors collected for that speaker. Thus, in Figure 1, the centroid of speaker A lies at the origin of the respective allophone vectors derived as speaker A uttered the following words: "mass", "lack", and "had". Thus the centroid for speaker A contains information that in rough terms represents the "average" phoneme 'ae' for that speaker.

[0017] By comparison, the centroid of speaker B lies to the right of speaker A in speaker space. Speaker B's centroid has been generated by the following utterances: "laugh", "rap," and "bag". As illustrated, the other speakers C-F lie in other regions within the speaker space. Note that each speaker has a set of allophones that are represented as vectors emanating from the centroid (three allophone vectors are illustrated in Figure 1). As illustrated, these vectors define angular relationships that are often roughly comparable between different speakers. Compare angle 10 of speaker A with angle 12 of speaker B. However, because the centroids of the respective speakers do not lie coincident with one another, the resulting allophones of speakers A and B are not the same. The present invention is designed to handle this problem by removing the speaker-dependent idiosyncrasies characterized by different centroid locations.

[0018] While the angular relationships among allophone vectors are generally comparable among speakers, that is not to say that the vectors are identical. Indeed, vector lengths may vary from one speaker to another. Male speakers and female speakers would likely have different allophone vector lengths from one another. Moreover, there can be

different angular relationships attributable to different speaker dialects. In this regard, compare angle 14 of speaker E with angle 10 of speaker A. This angular difference would be reflective, for example, where speaker A speaks a northern United States dialect whereas speaker E speaks a southern United States dialect.

[0019] These vector lengths and angular differences aside, the disparity in centroid locations represents a significant speaker-dependant artifact that conventional context dependent recognizers fail to address. As will be more fully explained below, the present invention provides a mechanism to readily compensate for the disparity in centroid locations and also to compensate for other vector length and angular differences.

[0020] Figure 2 illustrates a presently preferred first embodiment that we call the Eigen centroid plus delta tree embodiment. More specifically, Figure 2 shows the preferred steps for training the delta trees that are then used by the recognizer. Figures 3 and 4 then show alternate embodiments for use of that recognizer with speech supplied by a new speaker.

[0021] Referring to Figure 2, the delta decision trees used by this embodiment may be grown by providing acoustic data from a plurality of training speakers, as illustrated at 16. The acoustic data from each training speaker is projected or placed into an eigenspace 18. In the presently preferred embodiment the eigenspace can be truncated to reduce its size and computational complexity. We refer here to the reduced size eigenspace as K-space.

[0022] One procedure for creating eigenspace 18 is illustrated by steps 20-26. The procedure uses the training speaker acoustic data 16 to generate speaker dependent (SD) models for each training speaker, as depicted at step 20. These models are then vectorized at step 22. In the presently preferred embodiment, the speaker dependent models are vectorized by concatenating the parameters of the speech models of each speaker. Typically Hidden Markov Models are used, resulting in a supervector for each speaker that may comprise an ordered list of parameters (typically floating point numbers) corresponding to at least a portion of the parameters of the Hidden Markov Models for that speaker. The parameters may be organized in any convenient order. The order is not critical; however, once an order is adopted it must be followed for all training speakers.

[0023] Next, a dimensionality reduction step is performed on the supervectors at step 24 to define the eigenspace. Dimensionality reduction can be effected through any linear transformation that reduces the original high-dimensional supervectors into basis vectors. A non-exhaustive list of dimensionality reduction techniques includes: Principal Component Analysis (PCA), Independent Component Analysis (ICA), Linear Discriminate Analysis (LDA), Factor Analysis (FA) and Singular Value Decomposition (SVD).

[0024] The basis vectors generated at step 24 define an eigenspace spanned by the eigenvectors. Dimensionality reduction yields one eigenvector for each one of the training speakers. Thus if there are n training speakers, the dimensionality reduction step 24 produces n eigenvectors. These eigenvectors define what we call eigenvoice space or eigenspace.

[0025] The eigenvectors that make up the eigenspace each represent a different dimension across which different speakers may be differentiated. Each supervector in the original training set can be represented as a linear combination of these eigenvectors. The eigenvectors are ordered by their importance in modeling the data: the first eigenvector is more important than the second, which is more important than the third, and so on.

[0026] Although a maximum of n eigenvectors is produced at step 24, in practice, it is possible to discard several of these eigenvectors, keeping only the first K eigenvectors. Thus at step 26 we optionally extract K of the n eigenvectors to comprise a reduced parameter eigenspace or K-space. The higher order eigenvectors can be discarded because they typically contain less important information with which to discriminate among speakers. Reducing the eigenvoice space to fewer than the total number of training speakers helps to eliminate noise found in the original training data, and also provides an inherent data compression that can be helpful when constructing practical systems with limited memory and processor resources.

[0027] Having constructed the eigenspace 18, the acoustic data of each individual training speaker is projected or placed in eigenspace as at 28. The location of each speaker's data in eigenspace (K-space) represents each speaker's centroid or average phoneme pronunciation. As illustrated in Figure 1, these centroids may be expected to differ from speaker to speaker. Speed is one significant advantage of using the eigenspace technique to determine speaker phoneme centroids.

[0028] The presently preferred technique for placing each speaker's data within eigenspace involves a technique that we call the Maximum Likelihood Estimation Technique (MLED). In practical effect, the Maximum Likelihood Technique will select the supervector within eigenspace that is most consistent with the speaker's input speech, regardless of how much speech is actually available.

[0029] To illustrate, assume that the speaker is a young female native of Alabama. Upon receipt of a few uttered syllables from this speaker, the Maximum Likelihood Technique will select a point within eigenspace that represents all phonemes (even those not yet represented in the input speech) consistent with this speaker's native Alabama female accent.

[0030] The Maximum Likelihood Technique employs a probability function Q that represents the probability of generating the observed data for a predefined set of Hidden Markov Models. Manipulation of the probability function Q is

EP 1 103 952 B1

made easier if the function includes not only a probability term P but also the logarithm of that term, $\log P$. The probability function is then maximized by taking the derivative of the probability function individually with respect to each of the eigenvalues. For example, if the eigenspace is on dimension 100 this system calculates 100 derivatives of the probability function Q, setting each to zero and solving for the respective eigenvalue W.

5 [0031] The resulting set of Ws, so obtained, represents the eigenvalues needed to identify the point in eigenspace that corresponds to the point of maximum likelihood. Thus the set of Ws comprises a maximum likelihood vector in eigenspace. This maximum likelihood vector may then be used to construct a supervector that corresponds to the optimal point in eigenspace.

10 [0032] In the context of the maximum likelihood framework of the invention, we wish to maximize the likelihood of an observation O with regard to a given model. This may be done iteratively by maximizing the auxiliary function Q presented below.

15
$$Q(\lambda, \hat{\lambda}) = \sum_{\theta \in \text{states}} P(O, \theta | \lambda) \log [P(O, \theta | \hat{\lambda})]$$

where λ is the model and $\hat{\lambda}$ is the estimated model.

20 [0033] As a preliminary approximation, we might want to carry out a maximization with regards to the means only. In the context where the probability P is given by a set of HMMs, we obtain the following:

25
$$Q(\lambda, \hat{\lambda}) = \text{const} - \frac{1}{2} P(O | \lambda) \sum_{\substack{\text{states} \\ m \in \lambda}} \sum_{\substack{m \\ \text{mix} \\ \text{gauss}}} \sum_{\substack{t \\ \in S}} \left\{ \gamma_m^{(s)}(t) [n \log(2\pi) + \log |C_m^{(s)}| + h(o_t, m, s)] \right\}$$

where:

30
$$h(o_t, m, s) = (o_t - \hat{\mu}_m^{(s)})^T C_m^{(s)-1} (o_t - \hat{\mu}_m^{(s)})$$

and let:

35 o_t be the feature vector at time t
 $C_m^{(s)-1}$ be the inverse covariance for mixture Gaussian m of state s
 $\hat{\mu}_m^{(s)}$ be the approximated adapted mean for state s, mixture component m
 $\gamma_m^{(s)}(t)$ be the $P(\text{using mix Gaussian } m | \lambda, o_t)$

40 [0034] Suppose the Gaussian means for the HMMs of the new speaker are located in eigenspace. Let this space be spanned by the mean supervectors $\bar{\mu}_j$, with $j=1 \dots E$,

45
$$\bar{\mu}_j = \begin{bmatrix} \bar{\mu}_1^{(1)}(j) \\ \bar{\mu}_2^{(1)}(j) \\ \vdots \\ \bar{\mu}_m^{(1)}(j) \\ \vdots \\ \bar{\mu}_{M_s}^{(s)}(j) \end{bmatrix}$$

50 where $\bar{\mu}_m^{(s)}(j)$ represents the mean vector for the mixture Gaussian m in the state s of the eigenvector (eigenmodel) j.

[0035] Then we need:

5

$$\hat{\mu} = \sum_{j=1}^E w_j \bar{\mu}_j$$

[0036] The $\bar{\mu}_j$ are orthogonal and the w_j are the eigenvalues of our speaker model. We assume here that any new speaker can be modeled as a linear combination of our database of observed speakers. Then

10

$$\hat{\mu}_m^{(s)} = \sum_{j=1}^E w_j \bar{\mu}_m^{(s)}(j)$$

15

with s in states of λ , m in mixture Gaussians of M .

[0037] Since we need to maximize Q , we just need to set

20

$$\frac{\partial Q}{\partial w_e} = 0, \quad e=1..E.$$

(Note that because the eigenvectors are orthogonal, $\frac{\partial w_i}{\partial w_j} = 0, i \neq j$.)
Hence we have

25

$$\frac{\partial Q}{\partial w_e} = 0 = \sum_{\substack{\text{states} \\ m \\ \text{in } \lambda}} \sum_{\substack{\text{mixt} \\ \text{gauss} \\ t}} \sum_{\substack{\text{time} \\ mS \\ t}} \left\{ \frac{\partial}{\partial w_e} \gamma_m^{(s)}(t) h(o_e, s) \right\}, \quad e=1..E.$$

30

Computing the above derivative, we have:

35

$$0 = \sum_s \sum_m \sum_l \gamma_m^{(s)}(l) \left\{ -\bar{\mu}_m^{(s)T}(e) C_m^{(s)-1} o_e + \sum_{j=1}^E w_j \bar{\mu}_m^{(s)T}(j) C_m^{(s)-1} \bar{\mu}_m^{(s)}(e) \right\}$$

from which we find the set of linear equations

40

$$\sum_e \sum_s \sum_l \gamma_m^{(s)}(l) \bar{\mu}_m^{(s)T}(e) C_m^{(s)-1} o_e = \sum_s \sum_m \sum_l \gamma_m^{(s)}(l) \sum_{j=1}^E w_j \bar{\mu}_m^{(s)T}(j) C_m^{(s)-1} \bar{\mu}_m^{(s)}(e), \\ e=1..E.$$

45

[0038] Once the centroids for each speaker have been determined, they are subtracted at step 30 to yield speaker-adjusted acoustic data. Referring to Figure 1, this centroid subtraction process will tend to move all speakers within speaker space so that their centroids are coincident with one another. This, in effect, removes the speaker idiosyncrasies, leaving only the allophone-relevant data.

50

[0039] After all training speakers have been processed in this fashion, the resulting speaker-adjusted training data is used at step 32 to grow delta decision trees as illustrated diagrammatically at 34. A decision tree is grown in this fashion for each phoneme. The phoneme 'ae' is illustrated at 34. Each tree comprises a root node 36 containing a question about the context of the phoneme (i.e., a question about the phoneme's neighbors or other contextual information). The root node question may be answered either "yes" or "no", thereby branching left or right to a pair of child nodes. The child nodes can contain additional questions, as illustrated at 38, or a speech model, as illustrated at 40. Note that all leaf nodes (nodes 40, 42, and 44) contain speech models. These models are selected as being the models most suited for recognizing a particular allophone. Thus the speech models at the leaf nodes are context-dependent.

[0040] After the delta decision trees have been developed, as illustrated in Figure 1, the system may be used to recognize the speech of a new speaker. Two recognizer embodiments will now be described with reference to Figures 3 and 4. The recognizer embodiments differ essentially in whether the new speaker centroid is subtracted from the acoustic data prior to context-dependent recognition (Fig. 3); or whether the centroid information is added to the context-dependent models prior to context-dependent recognition.

[0041] Referring to Figure 3, the new speaker 50 supplies an utterance that is routed to several processing blocks, as illustrated. The utterance is supplied to a speaker-independent recognizer 52 that functions simply to initiate the MLED process.

[0042] Before the new speaker's utterance is submitted to the context-dependent recognizer 60, the new speaker's centroid information is subtracted from the speaker's acoustic data. This is accomplished by calculating the position of the new speaker within the eigenspace (K-space) as at 62 to thereby determine the centroid of the new speaker as at 64. Preferably the previously described MLED technique is used to calculate the position of the new speaker in K-space.

[0043] Having determined the centroid of the new speaker, the centroid data is subtracted from the new speaker's acoustic data as at 66. This yields speaker-adjusted acoustic data 68 that is then submitted to the context-dependent recognizer 60.

[0044] The alternate embodiment illustrated at Figure 4 works in a somewhat similar fashion. The new speaker's utterance is submitted to the speaker-independent recognizer 52 as before, to initiate the MLED process. Of course, if the MLED process is not being used in a particular embodiment, the speaker-independent recognizer may not be needed.

[0045] Meanwhile, the new speaker's utterance is placed into eigenspace as at step 62 to determine the centroid of the new speaker as at 64. The centroid information is then added to the context-dependent models as at 72 to yield a set of speaker-adjusted context-dependent models 74. These speaker-adjusted models are then used by the context-dependent recognizer 60 in producing the recognizer output 70. Table I below shows how exemplary data items for three speakers may be speaker-adjusted by subtracting out the centroid. All data items in the table are pronunciations of the phoneme 'ae' (in a variety of contexts).

[0046] Figure 5 then shows how a delta tree might be constructed using this speaker-adjusted data. Figure 6 then shows the grouping of the speaker-adjusted data in acoustic space. In Figure 6 +1 means next phoneme; the fricatives are the set of phonemes {f, h, s, th, ...}; voiced consonants are {b, d, g, ...}.

30

TABLE 1

Spkr1 : centroid = (2,3)					
	"half" = >	<h *ae f>	(3, 4)	- (2,3)	= (1,1)
	"sad" = >	<s *ae d>	(2, 2)	- (2,3)	= (0,-1)
	"fat" = >	<f *ae t>	(1.5, 3)	- (2,3)	= (-0.5, 0)
Spkr2 : centroid = (7,7)					
	"math" = >	< m *ae th>	(8,8)	- (7,7)	= (1,1)
	"babble" = >	< b *ae b l>	(7,6)	- (7,7)	= (0,-1)
	"gap" = >	< g *ae p>	(6.5, 7)	- (7,7)	= (-0.5,0)
Spkr3 : centroid = (10,2)					
	"task" = >	< t *ae s k>	(11,3)	- (10,2)	= (1,1)
	"cad" = >	< k *ae d>	(10,1)	- (10,2)	= (0,-1)
	"tap" = >	< t *ae p>	(9.5,2)	- (10,2)	= (-0.5,0)

50

[0047] If desired, standard deviations as well as means may be used in the speaker-adjustment process. This would be done, for example, by imposing a unit variance condition (as in cepstral normalization). After speaker-dependent centroid training, the supervectors submitted to MLED would contain standard deviations as well as means. For each training data item, after subtraction of the phoneme state centroid from each data item, the item would be further adjusted by dividing by centroid standard deviations. This would result in even more accurate pooling of allophone data by the trees. There would be some computational costs at run time when using this technique, because the speaker-adjustment of incoming frames would be slightly more complex.

[0048] As previously noted, co-articulation can be affected by speaker type in a way that causes the direction of the

allophone vectors to differ. This was illustrated in Figure 1 wherein the angular relationships of offset vectors differed depending on whether the speaker was from the north or from the south. This phenomenon may be taken into account by including decision tree questions about the eigen dimensions. Figure 7 shows an exemplary delta decision tree that includes questions about the eigen dimensions in determining which model to apply to a particular allophone. In Figure

7, questions 80 and 82 are eigen dimension questions. The questions ask whether a particular eigen dimension (in this case dimension 3) is greater than zero. Of course, other questions can also be asked about the eigen dimension.

[0049] Another embodiment of the invention will now be described in connection with Figure 8. This embodiment is suited for applications in which there is a sufficient quantity of data per speaker to train reasonably accurate speaker-dependent models. In this embodiment it is not necessary to ascertain the centroids of each speaker.

[0050] However, in order to employ the eigenvoice technique, it is necessary to have a set of supervectors (one from each training speaker-dependent model). These supervectors must have the same dimension and must be aligned in the same sense that the index i must represent the same parameter across all speaker-dependent models.

[0051] Therefore, to grow a good context-dependent allophone tree for a given phoneme that is sharable across speakers, this embodiment pools data across speakers, but keeps track of which data item came from which speaker.

[0052] The maximum likelihood estimation (MLE) criterion for choosing a question is thus extended to accumulate an overall score for each test question, while separately evaluating and retaining the scores for individual speakers. Figure 8 illustrates the technique.

[0053] Referring to Figure 8, the decision tree structure is grown by providing a pool of questions 100. These questions are individually tested by the tree-growing algorithm, to determine which questions best define the structure of the allophone trees.

[0054] Recall that each test question may relate in some way to the context in which a particular phoneme occurs. Thus the test question might be, for example, whether the given phoneme is preceded by a fricative. The tree-growing algorithm grows individual trees for each phoneme, starting with a root node question and proceeding to additional nodes, as needed until the allophones of that phoneme are well-represented by the tree structure. Selection of the root node question and any intermediate node questions proceeds as illustrated in Figure 8.

[0055] The procedure for selecting test questions works by assuming that the current question under evaluation (question 104) has been chosen for that node of the tree. Speaker data from the training speakers 106 are evaluated by the test question 104 to thereby split the speech data into two portions: a portion that answered "yes" to the test question and a portion that answered "no" to the test question. Speech models are then constructed using the test speaker data. Specifically, a "yes" model 106 and a "no" model 108 are constructed for each speaker. This is different from the conventional procedure in which data for all speakers is pooled and for a given question, one "yes" and one "no" model are trained from the pooled data. The models are trained by training acoustic features on all the speech data examples that answer "yes" to the test question and similarly training another set of acoustic features on the data that answers "no" to the test question.

[0056] After having generated a "yes" model 106 and a "no" model 108 for each speaker, the system calculates the probability score of all the "yes" data given the "yes" model 106 and also calculates the probability score of all the "no" data given the "no" model 108. A high probability score means that the constructed model is doing a good job of recognizing its portion of the training data. A low probability score means that the model, while perhaps the best model that could be created using the training data, is not doing a good job of recognizing the phoneme in question.

[0057] The probability scores are assessed to compute the overall score for the test question 104. The computation proceeds as illustrated in Figure 8 as follows. First the respective probability scores for the "yes" model and the "no" model are computed for a first training speaker (speaker A). These scores are multiplied together to give a cumulative product score indicative of how well the models worked with speaker A. This is illustrated at 112. The same procedure is then followed for the remaining training speakers, one speaker at a time, as illustrated at 114 and 116. Finally, when all of the training speakers have been taken into account, an overall score is computed by multiplying the resultant products derived from individual speakers. Thus the products ascertained at steps 112, 114 and 116 are multiplied together to yield an overall score for the test question at 118.

[0058] Having generated an overall score for the first test question, iterator 102 stores the overall score results and then draws a second question from the pool of questions 100 for testing in the same fashion. When all questions in the pool have been tested, the question that gave the best overall score is selected for that node of the decision tree.

[0059] After the root node of the decision tree has been determined as described above, the iterator 102 may continue to assess whether further intermediate nodes produce sufficient improvements in allophone recognition to warrant adding additional nodes to the tree. Ultimately, when the tree is grown in this fashion, the leaf nodes contain the models that best "recognize" the allophones of a particular phoneme.

[0060] After the decision tree structures have been identified through the foregoing procedures, the eigenvoice tech-

nique may now be applied. If a single Gaussian per leaf node is sufficient to represent the allophone, the allophonic speaker-dependent models are trained using the shared tree structure to obtain a set of supervectors which are then used to construct the eigenspace through dimensionality reduction. With training now complete, the online step is a simple MLED estimation of the eigenvoice coefficients. Multiple Gaussian models are slightly more complicated because the question of alignment must be addressed. That is, whereas it is known that leaf node N of speaker-dependent model 1 and leaf node N of speaker dependent model 2 represent the same allophone, it is not certain that Gaussian i of leaf N in speaker-dependent model 1 represents the same phenomenon as Gaussian i of leaf N in speaker dependent model 2. One way to address this issue is to find a centroid for each leaf of the speaker-dependent model and then speaker-adjust the data reaching all leaves. One would then pool data for a given leaf across speaker-dependent models and calculate shared delta Gaussians. At run time, MLED would yield estimates of all leaf centroids, which could then be subtracted from the new speaker's data before it is evaluated against the delta Gaussians.

Claims

15

1. A method for developing context-dependent models for automatic speech recognition, comprising:

generating an eigenspace (18) to represent a training speaker population;
providing a set of acoustic data (16) for at least one training speaker and representing said acoustic data in said eigenspace (28) to determine at least one allophone centroid for said training speaker; and
subtracting said centroid from said acoustic data (30) to generate speaker-adjusted acoustic data for said training speaker;
using said speaker-adjusted acoustic data to grow at least one decision tree (32) having leaf nodes containing context-dependent models for different allophones.

25

2. The method of claim 1 further comprising using a set of acoustic data for a plurality of training speakers to generate said speaker-adjusted acoustic data for each of said plurality of training speakers.

30

3. The method of claim 1 wherein said eigenspace is generated by constructing supervectors (22) based on speech from said training speaker population and performing dimensionality reduction (24) upon said supervectors to define a reduced dimensionality space that spans said training speaker population.

35

4. A method of performing speech recognition using said context-dependent models developed as recited in claim 1, comprising:

providing speech data from a new speaker (50);
using said eigenspace (62) to determine at least one new speaker centroid (64) of a new speaker and subtracting said new speaker centroid from said speech data (66) from said new speaker to generate speaker-adjusted data (68); and
applying said speaker-adjusted data to a speech recognizer (60) employing said context-dependent models (58).

45

5. A method of performing speech recognition using said context-dependent models developed as recited in claim 1, comprising;

providing speech data from a new speaker (50);
using said eigenspace (62) to determine at least one new speaker centroid (64) of a new speaker and adding said new speaker centroid to said context-dependent models (72) to generate new speaker-adjusted context-dependent models (74); and
applying said speech data to a speech recognizer (60) employing said new speaker-adjusted context-dependent models (74).

Patentansprüche

55

1. Verfahren zur Entwicklung kontextabhängiger Modelle zur automatischen Spracherkennung, aufweisend:

Erzeugen eines Eigenraums (18), um eine Schulungssprecherpopulation zu repräsentieren;

Bereitstellen einer Menge akustischer Daten (16) für mindestens einen Schulungssprecher und Darstellen der akustischen Daten im Eigenraum (28), um mindestens einen Allophon-Schwerpunkt für den Schulungssprecher zu bestimmen; und

Subtrahieren des Schwerpunkts von den akustischen Daten (30), um sprecherangepasste akustische Daten für den Schulungssprecher zu erzeugen;

Verwenden der sprecherangepassten akustischen Daten, um mindestens einen Entscheidungsbäum (32) mit Blattknoten zu züchten, die kontextabhängige Modelle für verschiedene Allophone haben.

2. Verfahren nach Anspruch 1, ferner die Verwendung einer Menge akustischer Daten für eine Mehrzahl Schulungssprecher aufweisend, um sprecherangepasste akustische Daten für jeden Sprecher der Mehrzahl Schulungssprecher zu erzeugen.

3. Verfahren nach Anspruch 1, bei dem der Eigenraum durch die Konstruktion von Supervektoren (22) auf Basis von Sprache der Schulungssprecherpopulation erzeugt und eine Dimensionalitätsverringerung (24) dieser Supervektoren ausgeführt wird, um einen verringerten Dimensionalitätsraum zu definieren, der die Schulungssprecherpopulation abdeckt.

4. Verfahren zur Ausführung von Spracherkennung unter Verwendung der kontextabhängigen Modelle, die gemäß Anspruch 1 entwickelt worden sind, aufweisend:

Bereitstellen von Sprachdaten von einem neuen Sprecher (50);

Verwenden des Eigenraums (62) zur Bestimmung mindestens eines neuen Sprecher-Schwerpunkts (64) und Subtrahieren des neuen Sprecher-Schwerpunkts von den Sprachdaten (66) des neuen Sprechers, um sprecherangepasste Daten (68) zu erzeugen; und

Liefern der sprecherangepassten Daten an einen Spracherkennner (60), in dem die kontextabhängigen Modelle (58) verwendet werden.

5. Verfahren zur Ausführung von Spracherkennung unter Verwendung der kontextabhängigen Modelle, die gemäß Anspruch 1 entwickelt worden sind, aufweisend:

Bereitstellen von Sprachdaten von einem neuen Sprecher (50);

Verwenden des Eigenraums (62) zur Bestimmung mindestens eines neuen Sprecher-Schwerpunkts (64) eines neuen Sprecher und Addieren des neuen Sprecher-Schwerpunkts zu den kontextabhängigen Modellen (72), um neue sprecherangepasste kontextabhängige Modelle (74) zu erzeugen; und

Liefern der Sprachdaten an einen Spracherkennner (60), in dem die neuen sprecherangepassten kontextabhängigen Modelle (74) verwendet werden.

Revendications

1. Procédé de développement de modèles dépendant du contexte pour la reconnaissance automatique de la parole, comprenant :

la production d'un espace propre (18) pour représenter une population de locuteurs d'apprentissage ;

la fourniture d'un ensemble de données acoustiques (16) pour au moins un locuteur d'apprentissage et la représentation des données acoustiques dans l'espace propre (28) de façon à déterminer au moins un centreïde d'allophones pour le locuteur d'apprentissage ; et

la soustraction du centreïde aux données acoustiques (30) pour produire des données acoustiques ajustées en fonction du locuteur pour le locuteur d'apprentissage ;

l'utilisation des données acoustiques ajustées en fonction du locuteur pour faire croître au moins un arbre de décision (32) ayant des noeuds d'extrémité contenant des modèles dépendant du contexte pour différents allophones.

2. Procédé suivant la revendication 1, comprenant en outre l'utilisation d'un ensemble de données acoustiques pour une pluralité de locuteurs d'apprentissage afin de produire des données acoustiques ajustées en fonction du locuteur pour chacun de la pluralité de locuteurs d'apprentissage.

3. Procédé suivant la revendication 1, dans lequel l'espace propre est produit en construisant des supervecteurs (22)

basés sur la parole provenant de la population de locuteurs d'apprentissage et en effectuant une réduction de dimensionnalité (24) sur ces supervecteurs pour définir un espace de dimensionnalité réduite couvrant la population de locuteurs d'apprentissage.

5 4. Procédé de reconnaissance de la parole utilisant les modèles dépendant du contexte développés à la revendication 1, comprenant :

10 la fourniture de données vocales par un nouveau locuteur (50) ;
l'utilisation de l'espace propre (62) pour déterminer au moins un centroïde de nouveau locuteur (64) et soustraire ce centroïde de nouveau locuteur aux données (66) vocales provenant du nouveau locuteur pour produire des données (68) ajustées en fonction du locuteur ; et
l'application des données ajustées en fonction du locuteur à un dispositif (60) de reconnaissance de la parole utilisant les modèles (58) dépendant du contexte.

15 5. Procédé pour effectuer la reconnaissance de la parole utilisant les modèles dépendant du contexte développés à la revendication 1, comprenant :

20 la fourniture de données vocales par un nouveau locuteur (50) ;
l'utilisation de l'espace propre (62) pour déterminer au moins un centroïde (64) de nouveau locuteur d'un nouveau locuteur et ajouter ce centroïde de nouveau locuteur aux modèles (72) dépendant du contexte pour produire des modèles (74) dépendant du contexte et ajustés en fonction du nouveau locuteur ; et
l'application des données vocales à un dispositif (60) de reconnaissance de la parole en utilisant les modèles (74) dépendant du contexte et ajustés en fonction du nouveau locuteur.

25

30

35

40

45

50

55





FIG. 1

FIG. 2






FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADING TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.